

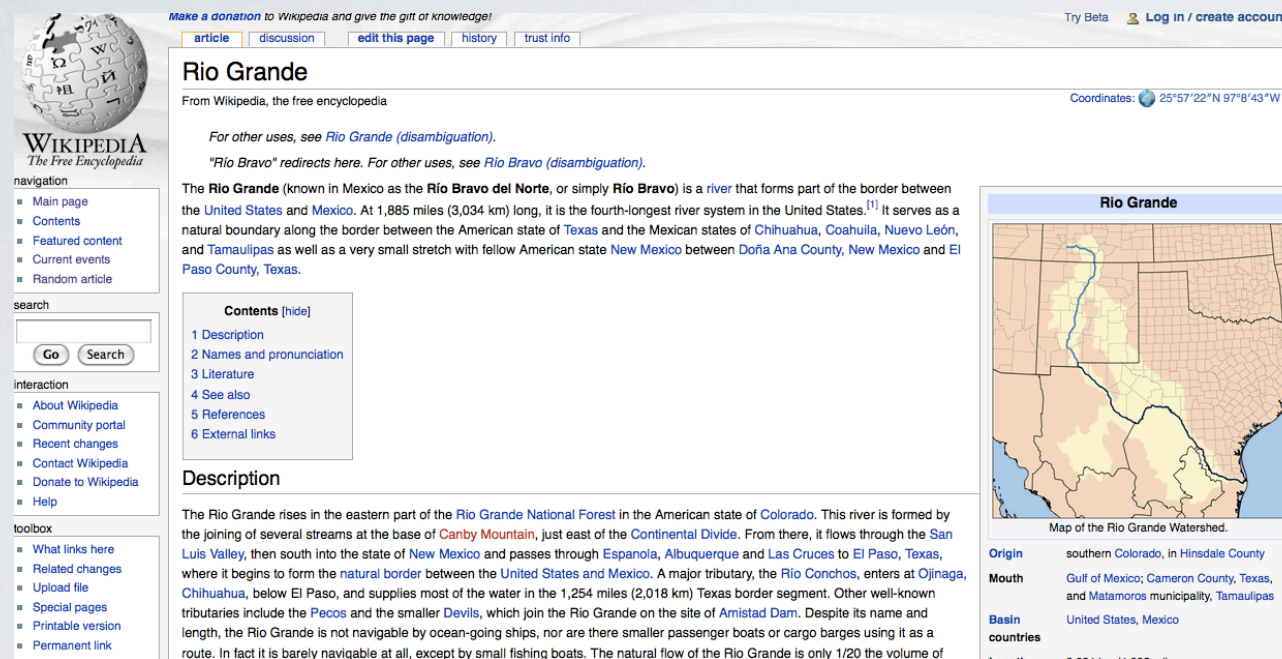


SCRAWL: A SEMANTIC CRAWLER FOR THE WIKIPEDIA AND BEYOND

Ian Pye, Luca de Alfaro
Shelly Spearing, Jorge Roman



THE PROBLEM



The Wikipedia is too big! We just want to look at the (potentially) interesting parts.

OUR SOLUTION

Target Page



Target Histogram



No

Close Enough?

Yes

Disk

Page to Evaluate

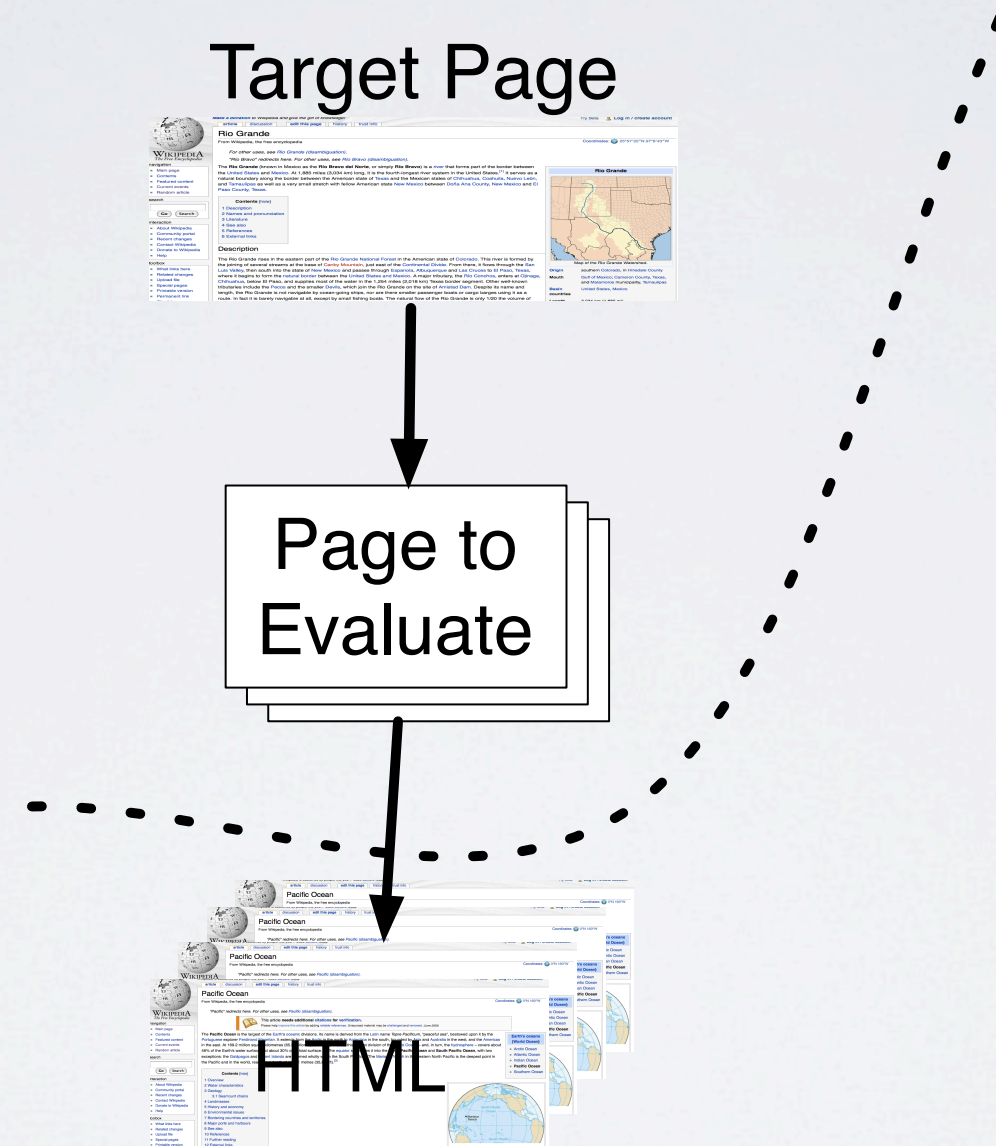
orch8.net,
extract
tri-grams

Revision
Histogram

HTML

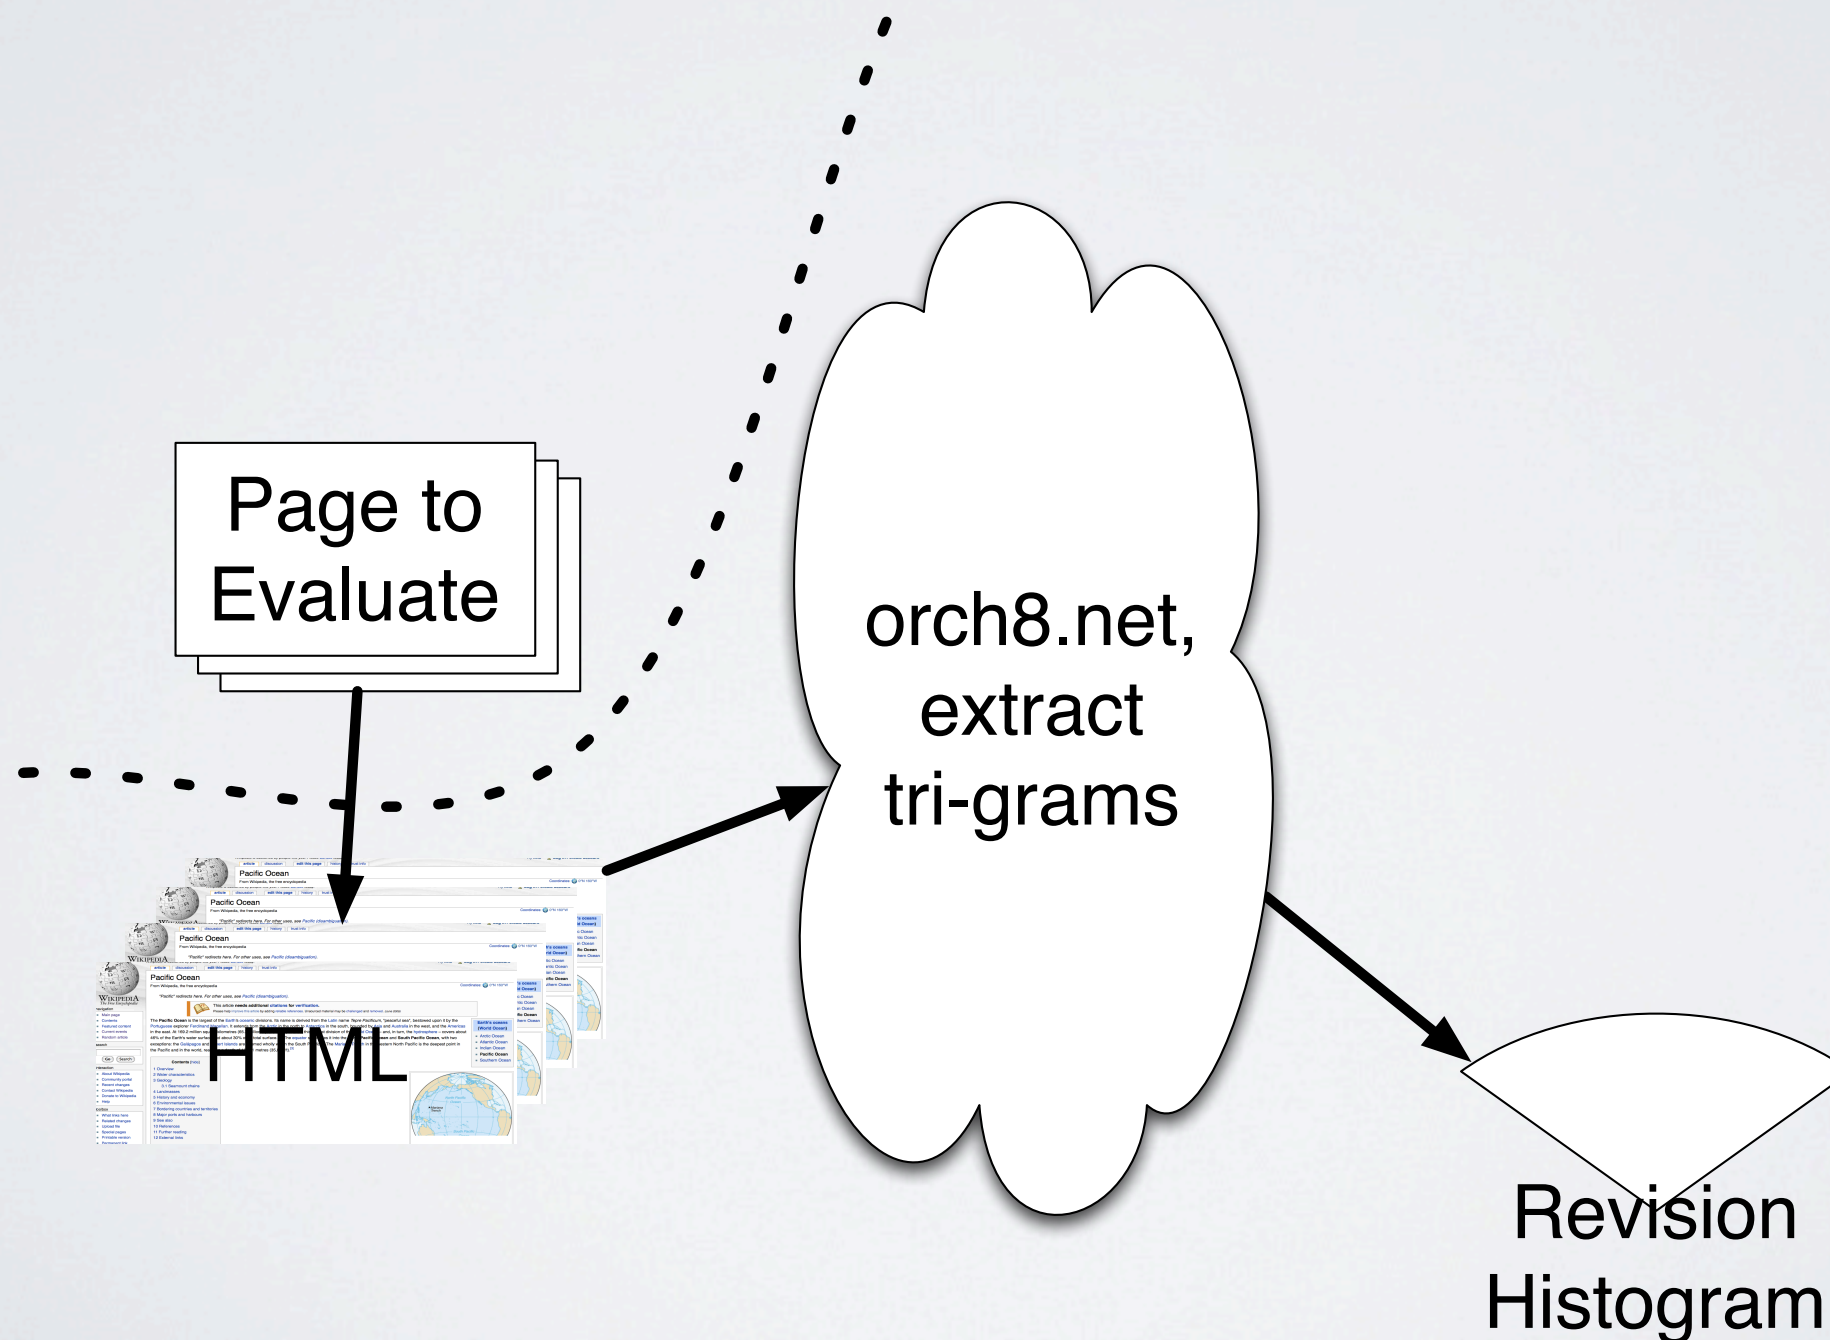
OUR SOLUTION

I) Download and Render to HTML




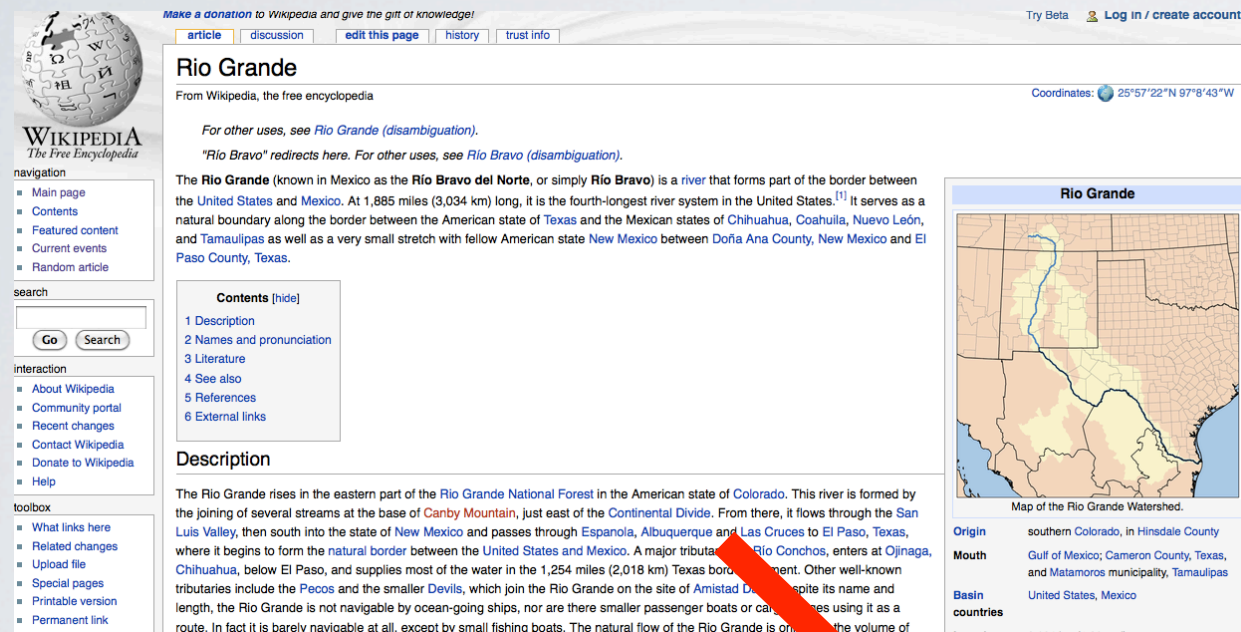
OUR SOLUTION

- 1) Download and Render to HTML
- 2) From HTML to Semantic Digest



OUR SOLUTION

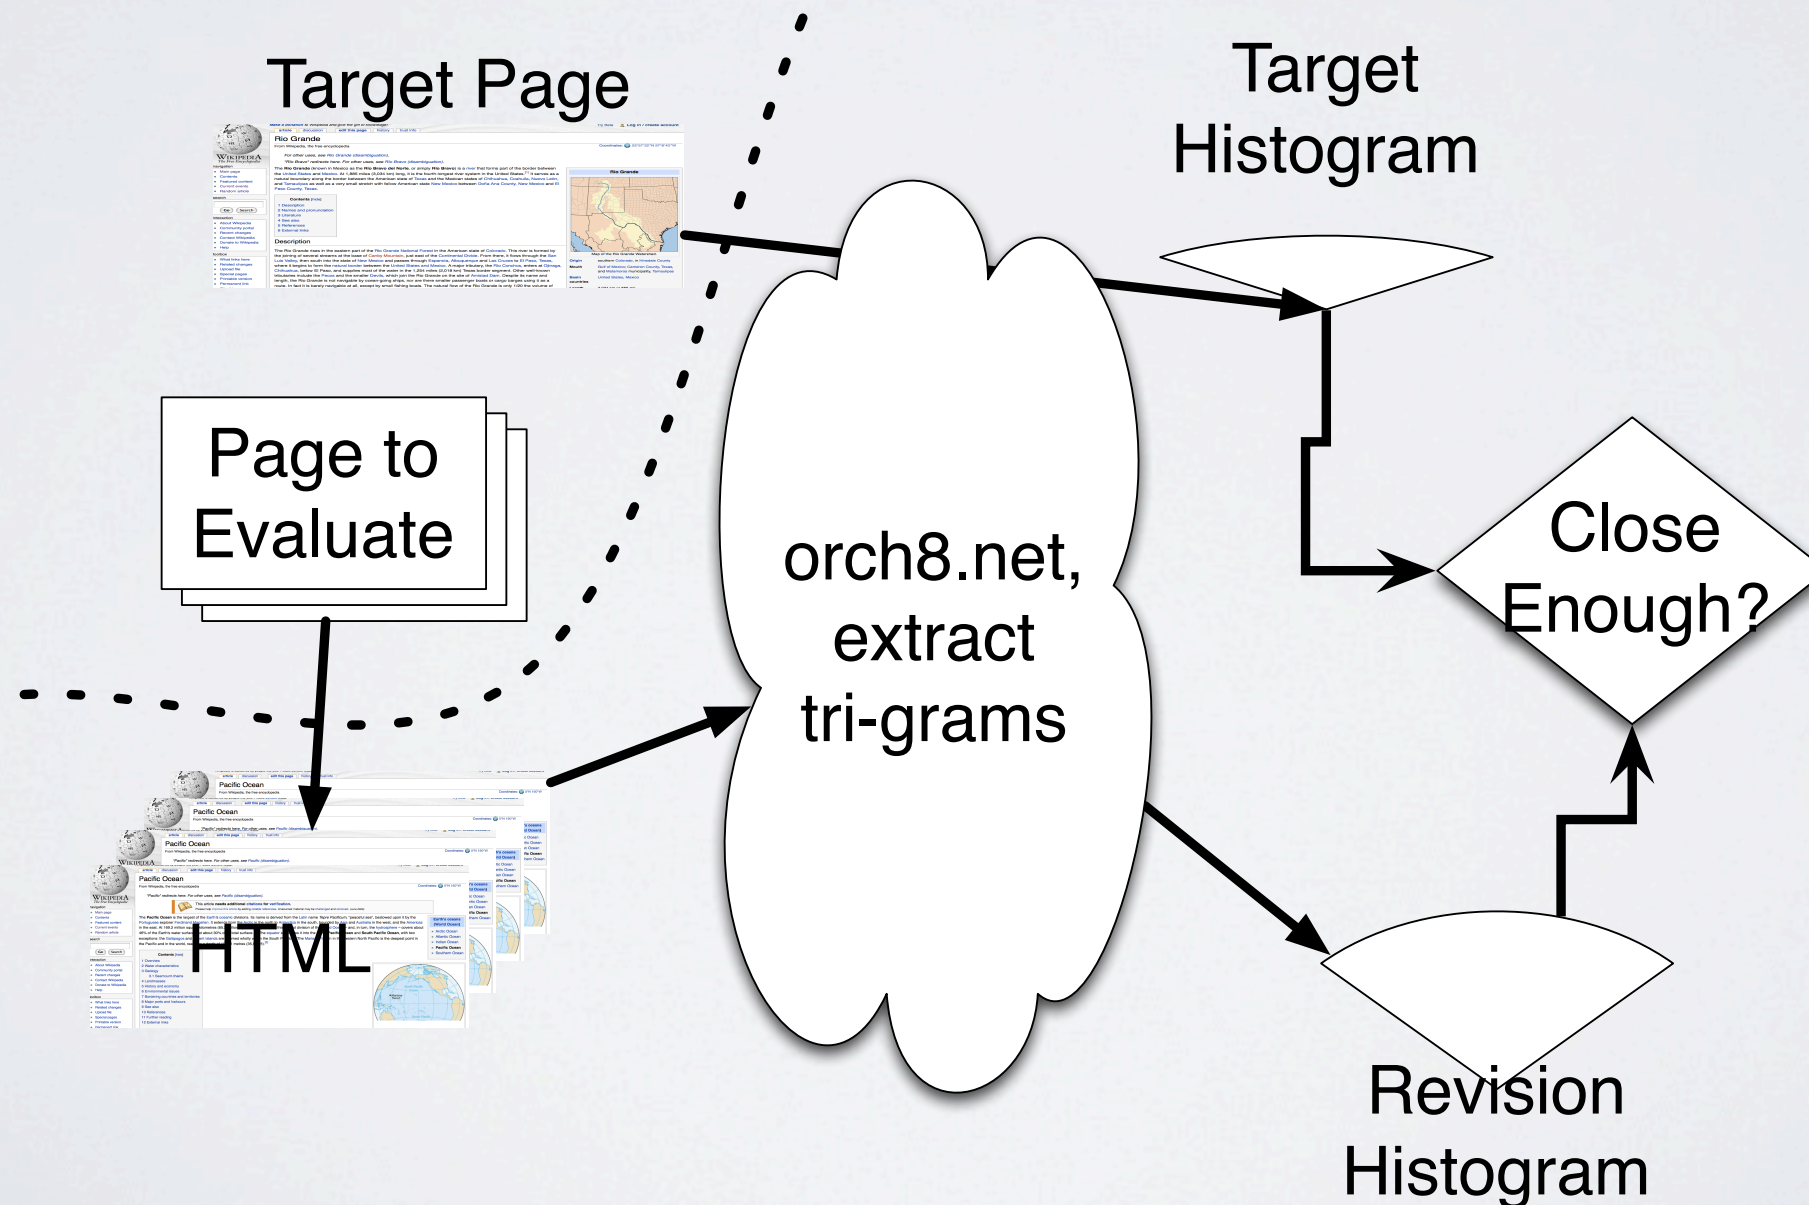
- 1) Download and Render to HTML
- 2) From HTML to Semantic Digest



Rio Grande River
Texas State River
Big Bend Park

OUR SOLUTION

- 1) Download and Render to HTML
- 2) From HTML to Semantic Digest
- 3) Euclidian Distance from the Target



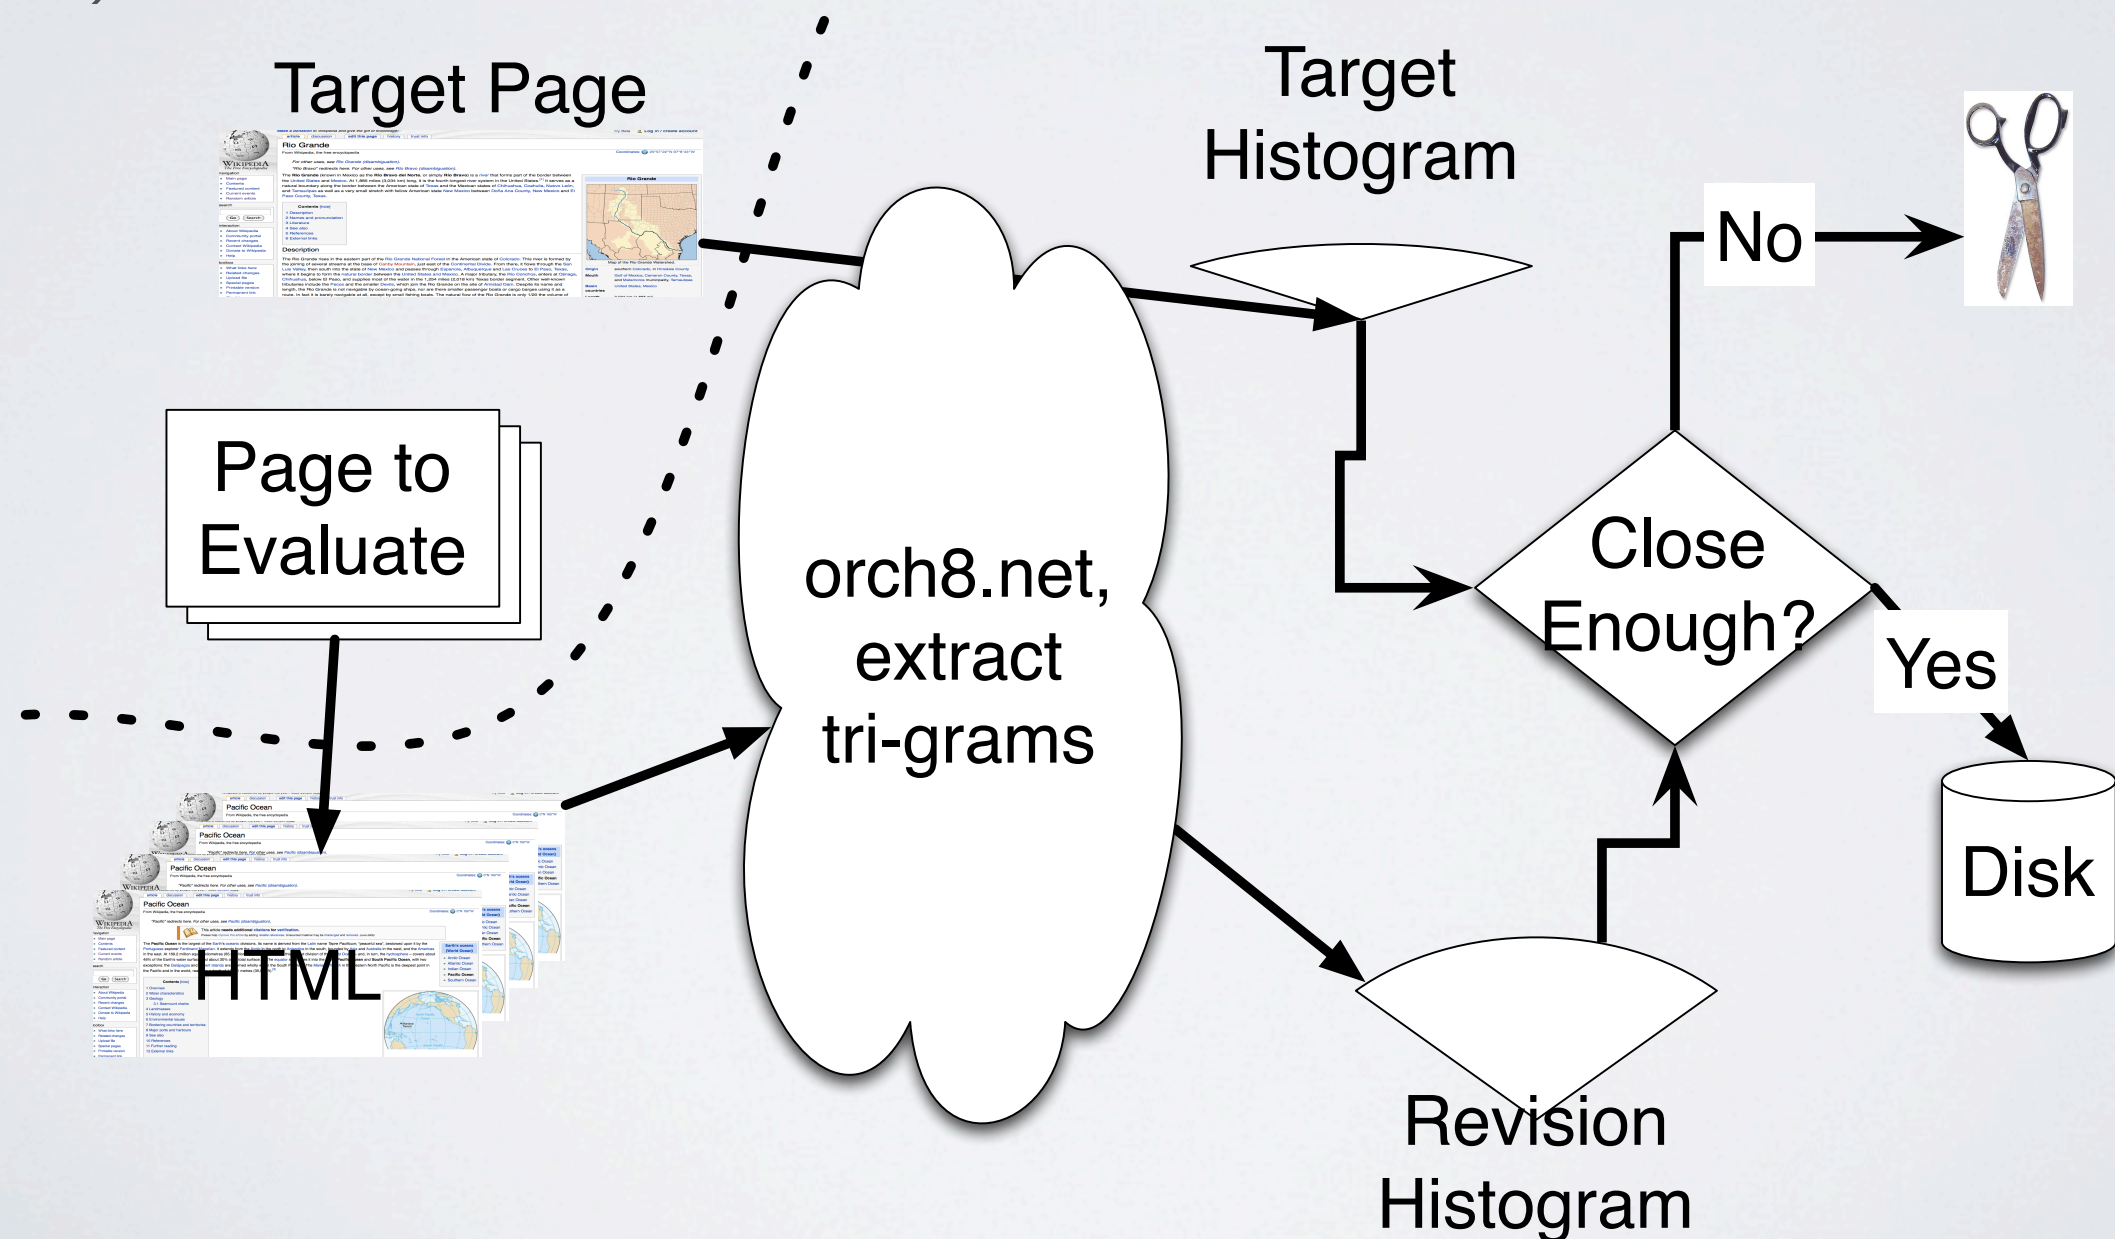
OUR SOLUTION

- 1) Download and Render to HTML
- 2) From HTML to Semantic Digest
- 3) Euclidian Distance from the Target



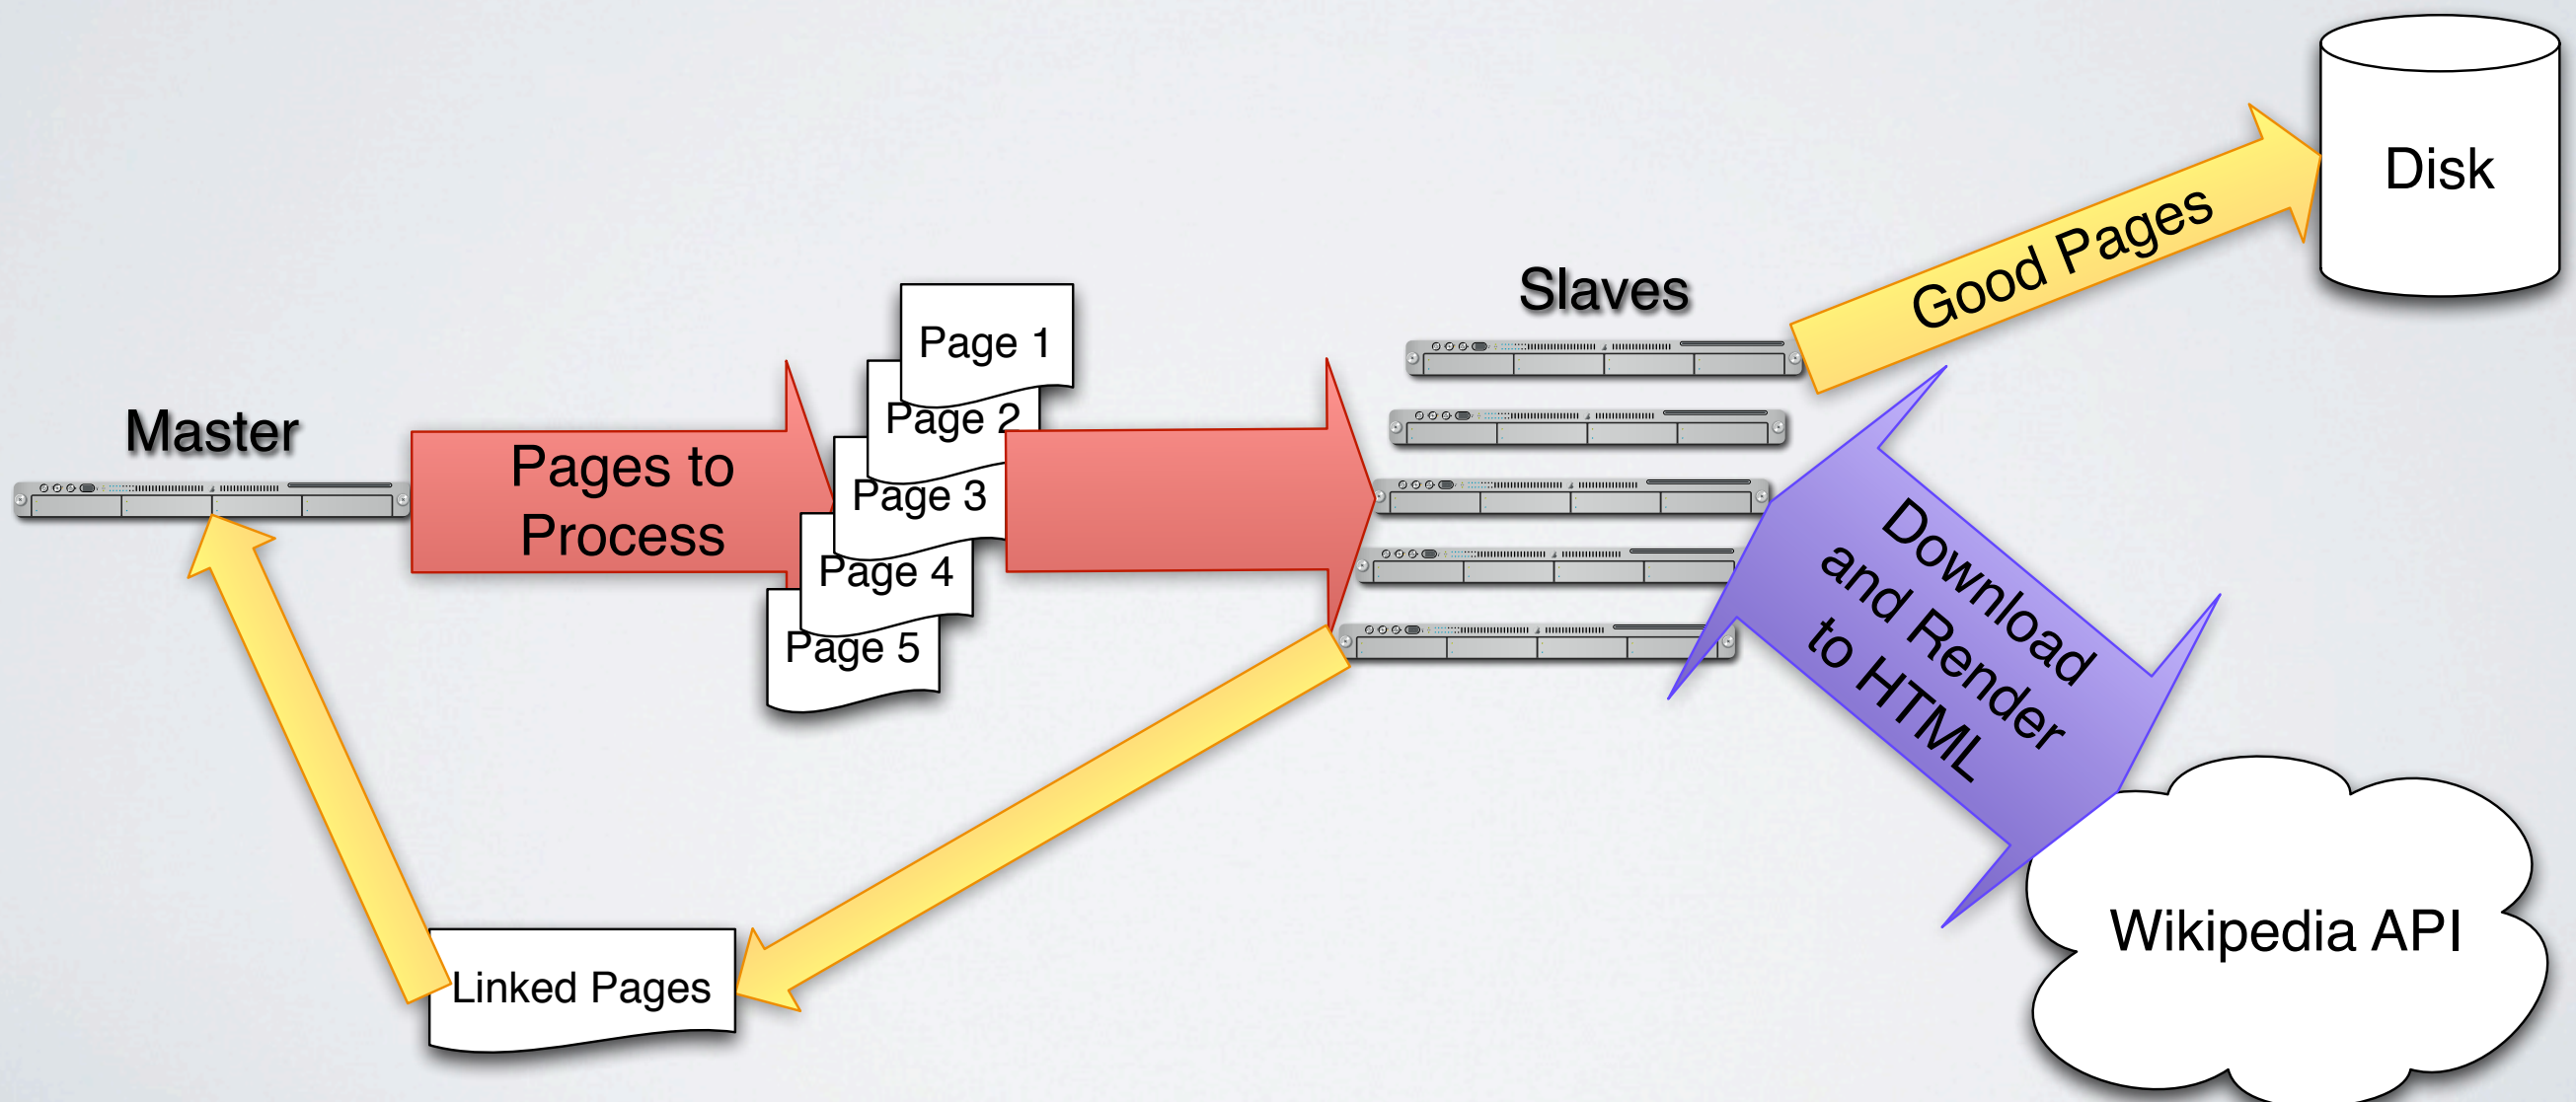
OUR SOLUTION

- 1) Download and Render to HTML
- 2) From HTML to Semantic Digest
- 3) Euclidian Distance from the Target
- 4) Extract Links and Recurse If Distance $<$ Max



DO THIS IN PARALLEL

OpenMPI allows us to run in a page-wise parallel fashion on a cluster.



TESTING PLATFORM: CRAIGZCRUZER [WWW.SUPERTRIC.COM](http://www.supertric.com)

The screenshot shows a Firefox browser window with the address bar set to <http://www.supertric.com/>. The page title is "CraigzCruzer". The navigation bar includes a link to "CraigzCruzer" and a user login status "ianpye@gmail.com" with a "logout" link. The main content area has a section "Pick a posting:" with a list item: <http://sfbay.craigslist.org/scz/apa/1418950362.html> | delete. Below this is a section "Add a new posting:" with two input fields. The first field is labeled "my posting url:" and has an example: (<http://sfbay.craigslist.org/scz/hou/1382798327.html>). The second field is labeled "my search url:" and has an example: (<http://sfbay.craigslist.org/search/apa/scz?query=&catAbbreviation=apa&minAsk=min&maxAsk=max&bedrooms=2&neighborhood=42>). Below the input fields is a "Sensitivity:" section with a slider set to 29 and a text box containing the number 29. A note below the slider states: "Higher values generate more results but also more false-positives." At the bottom of the form is a "Submit Query" button. The footer of the page contains the copyright notice "© 2009 SuperTric.com" and a link to the "README".

Firefox File Edit View History Bookmarks Tools Window Help

CraigzCruzer

<http://www.supertric.com/>

Most Visited OCA Bookmarklet Santa Cruz Public Li...

Gmail - Inbox... how_to_install... Google Docs -... Interesting W... F1 score - Wi... Massively coll...

[CraigzCruzer](#) | ianpye@gmail.com [logout](#) |

Pick a posting:

- <http://sfbay.craigslist.org/scz/apa/1418950362.html> | delete

Add a new posting:

my posting url:

example: (<http://sfbay.craigslist.org/scz/hou/1382798327.html>)

my search url:

example: (<http://sfbay.craigslist.org/search/apa/scz?query=&catAbbreviation=apa&minAsk=min&maxAsk=max&bedrooms=2&neighborhood=42>)

Sensitivity: 29

Higher values generate more results but also more false-positives.


© 2009 SuperTric.com | [README](#)

HOW ACCURATE IS IT?


METRICS (OR, A SHORT DIGRESSION INTO IR)

- Precision: the ratio of true positives to false positives
- Recall: the ratio of found positives to all positives
- F_β : How much to weight recall vs precision?

RESULTS: PRECISION AND RECALL



RESULTS: F_β



CONCLUSIONS

- SCrawl is a tool to extract what text MIGHT be interesting.
- Works though semantic meaning.
- Achieve high recall without sacrificing too much precision.
- Parallel processing for large datasets.